

Platform Solutions

Online News for Developers

Issue 22, July 30, 1999

Feature Story

Each month we provide a feature article on key industry trends and developments. Authored by a member of Intel's executive staff, the feature offers insightful information for product development, planning and execution.

IDF Preview

The Intel Developer Forum (IDF) conference is a semiannual, three-day conference targeted at business and consumer desktop, mobile, workstation, and server hardware manufacturers, and peripheral and software vendors working on advanced platform solutions. This special section previews just some of the 13 technology tracks comprising more than 100 in-depth sessions and 8 hands-on labs at the Fall '99 Conference coming up August 31 through September 2.

Inside Looking In

Senior Technical Marketing Manager for Platform Technologies Tim Mostad gives you a fresh perspective on the latest technologies making their way onto Intel Architecture platforms. Tim lets you see the work through the eyes of the people making it happen and lets you hear what they really think. It's straight talk from developers to developers.

Top Stories

Delivering in-depth reports on key platforms, products and technologies, our Top Stories provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Reader Services

If you're new to [Platform Solutions](#) and would like to receive this companion newsletter to the online version, please visit [Platform Solutions](#) on-line and go the "[Subscribe Now](#)" section to register and sign up for delivery. The online version provides direct links for quick access to the developer information and news reported in each issue, whether it's on Intel's Web site or industry Web sites.

We've created Platform Solutions to be your design news resource for the Intel® Architecture platform. If you can help us make it better, if you have a comment or a question or if you have a specific topic you would like to see covered, we'd like to hear from you. Just send an email to platform.solutions@intel.com.

If you do not want to receive this mailing in the future, please send an email to platform.solutions@intel.com with "unsubscribe" in the body of the message.

On behalf of all of us at Platform Solutions, welcome to the future of the PC platform!

DISCLAIMER: THE MATERIALS ARE PROVIDED "AS IS" WITHOUT ANY EXPRESS OR IMPLIED WARRANTY OF ANY KIND INCLUDING WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT OF INTELLECTUAL PROPERTY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT SHALL INTEL OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION, LOSS OF INFORMATION) ARISING OUT OF THE USE OF OR INABILITY TO USE THE MATERIALS, EVEN IF INTEL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME JURISDICTIONS PROHIBIT THE EXCLUSION OR LIMITATION OF LIABILITY FOR CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT APPLY TO YOU. INTEL FURTHER DOES NOT WARRANT THE ACCURACY OR COMPLETENESS OF THE INFORMATION, TEXT, GRAPHICS, LINKS OR OTHER ITEMS CONTAINED WITHIN THESE MATERIALS. INTEL MAY MAKE CHANGES TO THESE MATERIALS, OR TO THE PRODUCTS DESCRIBED THEREIN, AT ANY TIME WITHOUT NOTICE. INTEL MAKES NO COMMITMENT TO UPDATE THE MATERIALS.

Table of Contents

(Click on descriptions to jump to articles)

<i>Feature Story</i>	1
FEATURE	3
WIRED FOR MANAGEMENT SUMMIT '99 BRINGS INDUSTRY LEADERS TOGETHER ON AUGUST 2	3
IDF PREVIEW	5
LATEST IA-64 AND MERCED DETAILS AT FALL '99 IDF	5
ADDING EASE AND STYLE TO YOUR PC	8
DESIGNING NEXT-GENERATION PERFORMANCE DESKTOP SYSTEMS	11
EHOMES: INTERNET CONNECTIVITY THROUGHOUT THE HOME	13
INTERNET DRIVING CHANGE IN HIGH VOLUME SERVERS.....	16
APPLIED COMPUTING FIRMLY EMBEDDED AT IDF	18
INSIDE LOOKING IN	20
OF TECHNOLOGY AND FRUIT	20
TOP STORIES	22
DIRECT DRAM UPDATE: TEST RESULTS AVAILABLE ON THE WEB	22
NGIO SET TO OFFER NEXT-GENERATION I/O FOR THE INTERNET	24
FIVE STEPS TO SCALABLE LINUX APPLICATIONS	27
INCREASE DATA PROTECTION WITH THE HARDWARE-BASED INTEL® RANDOM NUMBER GENERATOR	30
DESIGNING FOR HIGH SERVER AVAILABILITY	32
DIGITAL VISUAL INTERFACE COMES INTO FOCUS	35
INTEL® MOTHERBOARDS: RELIABILITY AND SUPPORT PROVIDE SYSTEM INTEGRATORS WITH A COMPETITIVE EDGE	39
TECHNOLOGY NEWS BYTES	42
<i>Bluetooth(tm) Special Interest Group Announces Release of the Bluetooth 1.0 Specification</i>	42
<i>NGIO Forum Announces Release of 1.0 Specification</i>	42
<i>Intel Introduces File Sharing Appliance for Small Business Customers</i>	42
<i>Intel Enters Broadband Market to Enable Faster Internet Access for PC Users</i>	42
<i>Intel Forms a New Business Unit and Aligns Existing Groups around Internet Opportunities</i>	42
<i>Intel Completes Merger with Dialogic</i>	42
<i>Intel to Acquire Softcom</i>	42
<i>Leading Corporations Move to Intel and Windows NT-Based Workstations</i>	43
<i>Intel and SGI Announce Agreement to Optimize OpenGL API for Current and Future Intel Architecture Platforms</i>	43
<i>Intel Launches Worldwide Channel Initiative for Internet Service Providers</i>	43
INDUSTRY EVENTS	44
INDUSTRY EVENTS FOR AUGUST.....	44
INDUSTRY EVENTS FOR SEPTEMBER	46
INDUSTRY EVENTS FOR OCTOBER	47

Feature

Wired for Management Summit '99 Brings Industry Leaders Together on August 2

Todd Harrison
WfM Product Marketing Engineer
Intel Corporation

In response to the growing demand for information about Wired for Management-based products and technologies, Intel is hosting WfM Summit '99. Find out what's expected when CIOs and IT professionals convene and meet with leaders in the world of WfM-enabled solutions.

WfM 2.0 enhances the four basic management capabilities addressed in its predecessor WfM 1.1a, (system component instrumentation, universal network boot, power management, and remote wake-up), by implementing several changes. These three are key:

- Power management requirements have changed from APM to ACPI (using system states 1-5).
- The Universal Network Boot now enables boot server arbitration.
- SDKs have been added for streamlined help desk support capabilities and Boot Integrity Services.

In response to the growing demand for information about WfM-based products and technologies, Intel is hosting WfM Summit '99, a one-day conference to be held on August 2 at the Santa Clara Convention Center in the heart of Silicon Valley. Some 500 CIOs, business decision makers, and IT managers and directors are expected to attend the conference, where they'll have the opportunity to learn about the latest WfM 2.0-enabled offerings from PC manufacturers and management application vendors. In addition, attendees will be able to review case studies focusing on the use of WfM-enabled technologies at a number of leading enterprise sites.

Keynote Address by Gartner Group's Knox

WfM Summit '99 kicks off at 12:30 PM on August 2 with a keynote address by Kevin Knox, a leading total cost of ownership (TCO) analyst at the Gartner Group, one of the industry's most prominent market research firms. Knox will cover topics including technology planning and deployment, physical management, and support strategies, related technological trends, and other TCO issues currently affecting business enterprises.

Following Knox's keynote address, Intel's L.D. Weller—the emcee of the event and the current chairman of the Distributed Management Task Force (DMTF)—will discuss the benefits of implementing WfM-enabled technologies and present several case studies that examine how companies such as Hewitt Associates and AutoNation are using WfM-enabled solutions to reduce costs, increase efficiency, and enhance their business flexibility. The case studies feature video clips of IT managers who are already leveraging the advantages of WfM-enabled solutions across their enterprise computing environments.

WfM Vendors Showcase Their Wares

The widespread adoption of Wired for Management has paved the way for the introduction of numerous WfM-enabled solutions from system manufacturers and management application vendors alike. Following the presentations by the Gartner Group's Knox and Intel's Weller, WfM Summit attendees will have the opportunity to learn about the latest solutions from a score of vendors that are scheduled to provide Demo Showcases at the conference.

Eleven of today's PC manufacturers—Acer, Compaq, Dell, Fujitsu, Gateway, Hewlett-Packard, IBM, Micron, Quantec, Packard Bell/NEC and Toshiba—will be demonstrating WfM-enabled systems at the Convention Center. In addition, nine management application vendors will be on hand to demonstrate the latest enhancements to their WfM application solutions. These vendors include Altiris, Computer Associates, Hewlett-Packard, Intel, Microsoft, On Technology, Sterling Commerce, Tivoli, and Xpoint.

The Showcase will enable IT professionals and decision-makers to mingle with the industry's leading PC management solutions providers and learn more about how to leverage the advantages that WfM provides in their own enterprise computing environments.

Register Today to Attend the Free WfM Summit

WfM Summit '99 provides CIOs and IT professionals with a chance to convene at one location and meet many of the leading companies and individuals in the world of WfM-enabled solutions. There is no charge to attend the WfM Summit '99; however, registration is required. The registration Web page is located at [Intel's Wired for Management Summit '99](#). In addition, attendees will have the opportunity to win a laptop or desktop PC in a drawing that will be held at the conference.

The WfM initiative serves as an “umbrella” for a number of technologies that, when combined in a PC, provide a strong baseline of manageability features. WfM Summit '99 paves the way for implementing these features in today's enterprise computing environments and represents another in Intel's enabling efforts to support key technology-based initiatives designed to move the frontiers of computing and communications forward.

About the Author

Todd Harrison is Intel's product marketing engineer for the Wired for Management initiative, where his responsibilities include managing a wide range of WfM-related technical, marketing, and industry-enabling activities. He has worked at Intel for 8 years, during which time he has been involved in many different aspects of manageability technology.

Todd spent several years as a technical support engineer for the Intel® LANDesk® products and most recently served as a technical marketing engineer focused on the WfM initiative. He holds a Bachelor of Science degree in Electronic Engineering Technology from Weber State University.

For More Information

For more details and information on Wired for Management and the upcoming WfM Summit '99, please visit the following Web sites:

[WfM Summit '99](#)

[WfM Technology](#)

[WfM from an IT perspective](#)

IDF Preview

The Intel Developer Forum (IDF) conference is a semiannual, three-day conference targeted at business and consumer desktop, mobile, workstation, and server hardware manufacturers, and peripheral and software vendors working on advanced platform solutions. This special section previews just some of the 13 technology tracks comprising more than 100 in-depth sessions and 8 hands-on labs at the Fall '99 Conference coming up August 31 through September 2..

Latest IA-64 and Merced Details at Fall '99 IDF
Adding Ease and Style to your PC
Designing Next-Generation Performance Desktop Systems
eHomes: Internet Connectivity Throughout the Home
Internet Driving Change in High Volume Servers
Applied Computing Firmly Embedded at IDF

Latest IA-64 and Merced Details at Fall '99 IDF

Ron Curry
Director of Marketing
IA-64 Processor Division
Intel Corporation

To extend the Intel® Architecture further into the high-end server and workstation market segments, Intel has created the IA-64 architecture. IA-64 meets the increasing requirements of high end E-Business, decision support, and technical computing applications with higher levels of performance, scalability, availability, and headroom for future growth. The IA-64 architecture has been designed from the ground up to take advantage of recent advances in processor and compiler technology and deliver a strong roadmap that includes IA-32 instruction set compatibility. By overcoming performance limitations in traditional architectures, IA-64 provides developers an opportunity to build high-performance server and workstation solutions using common building blocks. This allows developers to leverage their development investment across multiple operating environments and deliver superior value to end-users.

Intel's IA-64 architecture—whose first product, the Merced processor, is scheduled for production in mid-2000—has been designed for high performance computing. Utilizing innovative techniques such as predication, speculation, and advanced register rotation allows IA-64 processors to remove hard-to-predict code branches, minimize the effect of memory latency, and reduce procedure call and loop processing overhead. These techniques combined with allowing the compiler to take a greater role in scheduling instruction execution, results in an architecture that gets more work done in each clock cycle. These features also enable greater headroom for future IA-64 processor performance growth by allowing more room on the processor for execution units and data caches.

Intel has created a comprehensive enabling program to ensure a variety of solution choices from development tools, to multiple operating systems, to compelling application stacks for servers and workstations. Vendors interested in working with the IA-64 architecture can now draw upon a wealth of information to support their product development efforts:

- A \$250 million investment fund, the Intel 64 Fund, has been established to invest in emerging technologies for IA-64
- An IA-64 developers' web site has been set up to give developers IA-64 documentation and links to tools

- The IA-64 Application Developers' Architecture Guide (ADAG) was publicly released to detail the IA-64 application environment
- At the upcoming Intel Developer Forum (IDF) Fall '99 in Palm Springs, California from August 31-September 2, Intel will deliver an IA-64 Track of 14 presentations to further aid hardware and software developers who haven't already started on IA-64 in creating solutions for the IA-64 architecture

The Intel 64 Fund and IA-64 ADAG

Vendor commitment to the IA-64 architecture has been unprecedented. Almost 20 leading server and workstation OEMs are developing IA-64 systems, seven of the major high-end operating systems are running on the Merced processor simulator, and dozens of leading server and workstation ISVs are actively porting their applications to IA-64. This commitment gives further proof of the benefits IA-64 provides, and ensures end users that they'll have access to complete solutions from the best-in-class vendors without being locked into a proprietary model.

To further broaden IA-64 development efforts, Intel has created the Intel 64 Fund and publicly released the IA-64 Application Developers' Architecture Guide (ADAG):

- Launched in May 1999, the *Intel 64 Fund* has raised more than \$250 million for the IA-64 development community. The fund consists of two components. The first component is a fund set up by Intel and OEMs Compaq, Dell, Hewlett-Packard, NEC and SGI to invest in companies developing tools and technologies for IA-64-based Internet, enterprise, and workstation applications. The second component is comprised of leading information technology (IT)-based companies such as Bank of America, Circuit City, Ford, General Electric and Morgan Stanley Dean Witter. These IT leaders will provide guidance to the companies the fund invests in to ensure solutions that better meet end-user needs. Developers interested in applying for funding should visit the [Intel 64 Fund Web site](#).
- The IA-64 ADAG (Application Developer's Architecture Guide) is a document that provides detailed descriptions of the IA-64 application environment including each instruction, its parameters, and valuable information on performance optimization tuning. This document gives developers the information they need to get the most from their applications on IA-64. Published on May 26, the IA-64 ADAG can be purchased in printed form by calling 1-800-548-4725 and referencing Document Order #245188, or downloaded for free as an [Adobe Acrobat PDF file](#).

IA-64 Architecture at IDF Fall '99

Information on the IA-64 industry effort will be a major focal point of the Intel Developer Forum (IDF) Fall '99, slated for August 31-September 2 at the Palm Springs Convention Center in Palm Springs, California. A special IA-64 track will provide developers with a wealth of pertinent information on products, resources, and timelines. Taught by technical experts from Intel and the industry, here's a closer look at some of the presentations:

- The **IA-64 Solutions Overview** is a one-hour discussion of the tremendous opportunity the IA-64 architecture provides developers. The overview will focus on the IA-64 roadmap, and how the performance and features of the Merced processor are poised to dramatically redefine the high-end server and workstation landscape. The course will be presented by the article author, Ronald Curry, director of marketing in Intel's IA-64 Processor Division.
- At the one-hour **IA 64 Software Program Overview**, Intel director of enterprise programs Mike Pope will discuss the IA-64 value proposition for server and workstation applications. He will focus on Intel's IA-64 software enabling programs and the progress being made by ISVs that have already started on IA-64. He will also detail key next steps for software developers who want to get started on IA-64 development. Additional presentations will describe the IA-64 compiler technology, specifics on getting software ready for IA-64, and development tools for each IA-64 operating system.

- The one-hour **Merced System Overview** describes the Intel Merced platform offerings for OEMs. Intel senior system architect, Bassam Elkhoury, and IA-64 lead technical marketing engineer, Mark Swearingen, will describe for the first time publicly, the high end features and capabilities of Intel's server and workstation platforms. Additional presentations will focus on the IA-64 Developers Interface Guide (DIG64) and the Extensible Firmware Interface, Intel's new system boot architecture.
- **Understanding the IA-64 Instruction Set Architecture (ISA)** is a two-hour course led by Intel senior architect, Gautam Doshi, which will focus on performance bottlenecks in current architectures, IA-64 features that address these bottlenecks, high-end application characteristics, and how IA-64 features benefit these characteristics.

For more information on these courses, and how to register to attend IDF Fall '99, please visit the [IDF Web site](#).

The Time for Action Is Now

The growth of the Internet and the creation of a significant new architecture represents an inflection point in which lies an opportunity for developers. Companies who take a leadership role have the best chance to benefit. With the Merced processor targeted for production around mid-2000, developers need to get started on their plans for IA-64 today. Many of the industry leading server and workstation [hardware and software vendors](#) are already well underway in creating their Merced processor based solutions.

Through programs such as the Intel 64 Fund, documentation and information sources such as the IA-64 ADAG and Intel's Developer Web site, and IA-64-focused learning opportunities such as IDF Fall '99, Intel is working to enable broader IA-64 development. The message is clear: hardware and software developers that have not yet begun IA-64 product development efforts should begin now. IA-64 represents a substantial opportunity, and the tools, education, and developer programs are in place for you to get started today.

About the Author

Ronald E. Curry is the director of marketing for the IA-64 Processor Division at Intel. He is responsible for product marketing, technical support, and industry enabling for IA-64 microprocessor products. Mr. Curry joined Intel in 1983 and has held several field and managerial positions including regional applications specialist, product marketing manager for the Intel® Pentium® processor, product marketing manager for IA-64 processors and chipsets, and marketing manager for Intel's Performance Marketing program.

For More Information

- [IA-64 Architecture Home Page](#)
- [The Intel 64 Fund](#)
- [The IA-64 Application Developers' Architecture Guide \(ADAG\)](#)
- [OEM IA-64 Development Efforts](#)
- [The DIG64 Developer's Interface Guide](#)
- [IA-64 at IDF Fall '99](#)

Adding Ease and Style to Your PC

Nancy Sumrall
PC EOU Marketing
Desktop Products Group
Intel Corporation

We've come a long way since the first Concept PC was shown at the September 1998 Intel Developer Forum. Intel introduced the first Legacy Removal Concept PC. Housed in a futuristic, pyramid chassis, the "Aztec," or "Yaquina" as it is now known, became one of the most talked about announcements at the Conference. The Yaquina (Figure 1) has traveled the world, and it has been shown and photographed all over the US, Europe, and Asia.



Figure1. The Yaquina Concept PC

The February 1999 IDF showcased at least eight new Concept PC designs. They included the Tetra, with a raspberry-colored chassis that weighed in around 4 lbs. and featured a CD-ROM visible from the front of the chassis, and the "Silicon Bonsai," a PC that grows upward, as well as several others. You can view them at the [IDF News Bureau PC Fashion Show](#).

A joint Easy PC Initiative between Intel and Microsoft was launched in April of this year to accelerate delivery of easier to use consumer PCs. Intel is providing specific building blocks to improve ease of use such as new chipsets, Hardware Implementation Guides, and the FlexATX motherboard. Intel is driving technology initiatives like USB 2.0, Instantly Available, Legacy Removal, and Form Factor to make the PC easier to expand, operate, and maintain—and look good, too.

These building blocks and technology initiatives are the basis for the Ease of Use Track at this Fall's IDF. Within the Ease of Use Track you'll find plenty of information on all the building blocks and technologies.

New chipsets now enable Instantly Available PCs. Instantly Available PC technology has been on the industry landscape for several years now. Initial designs were based on the Intel® 440BX chipset. As the industry shifts to Intel's next-generation chipset, the Intel® 810, new design considerations are required in order to support Instantly Available PC technology. This session tackles all the system design considerations, from board layout, peripheral requirements, and system BIOS needs. This session also gives insight into proper application design: what the software developer needs to know in order to take advantage of the new platform capabilities of the Instantly Available PC.

The Ease of Use PC Hardware Implementation Guide (HWIG '99) Rev. 0.9 was announced and presented at the February IDF. The HWIG '00 will be reviewed this Fall. The HWIG session provides an understanding of the **Easier to Use Hardware Implementation Guide '00 Rev 1.0**. It will focus on Legacy Removal and Instantly Available features and the relationships between innovative form factors and the Easy PC Initiative as well as other ease-of-use technology and HW/SW building blocks. An attendee in this session is sure to walk away with a plan for Easy PC products for Back to School '00.

The Ease of Use track will offer a session on "Implementing Easier to Use PC Designs using Standard Building Blocks." This session examines how to use basic system building blocks to achieve an Easy PC design. A decision matrix based on the basic system building blocks is reviewed to guide system/product designers through trade-offs. A **FlexATX** motherboard and chassis are used to demonstrate these concepts in a sample system everyone will be able to examine. It will also take a look at current form factor market trends toward smaller designs.

Intel and other industry leaders developed USB a few years ago, and it's now gaining significant momentum in the marketplace. All new PCs have USB ports, and most peripheral vendors are designing peripherals around USB. In certain device categories, conferencing cameras, for example, USB has a majority share of the volume sold. USB is well-known for its ease of use and faster performance than serial or parallel ports. At the Feb'99 IDF, Intel and six other "promoter" companies announced development of USB 2.0, the next generation of the spec that delivers at least 20 times the speed of USB. The USB 2.0 Technology session will cover **progress on the new USB 2.0 specification** (including an update on the speed) and will set the stage for the USB 2.0 Developers Conference in October.

One of the most popular Ease of Use sessions at the last IDF was "Ease of Use Benchmarks." The Fall IDF session will cover the elements and benchmarks key to providing easier to use PCs, focusing on the **Intel® Initial Experience Predictor (IEP) tool** and how this specific benchmark was built. Attendees will also learn how to use this tool as it applies to easier PC setup as well as how to score a system with IEP and how to improve their IEP score.

The Operating System is a major component of Ease of Use. Microsoft will offer a **Consumer Windows* OS** session. Microsoft will present and provide demonstrations of the Windows 98* Second Edition Operating System (OS) and focus on the further ease of use attributes that are planned for the next consumer release of the Windows OS due to ship in 2000. Ease of use aspects of the "Out of Box Experience" and day-to-day operations of the platform will be examined. Microsoft will provide an industry call to action on what steps PC OEMs, IHVs, and ISVs need to be taking to prepare the platform for these products and ensure a smooth transition for the consumer.

Several new Concept PC designs are in the works as are Easier to Use PCs that will be available later this year. What do the new Concept PCs look like? Bring your camera to the Fall '99 IDF Conference. We'll be demonstrating and showing you the latest in Palm Springs, August 31 through September 2.

About the Author

Nancy Sumrall is in PC EOU Marketing for Intel's Desktop Products Group, where her responsibilities focus on the development of industry initiatives to promote PC ease of use, technology, and conceptual design.

For More Information

[Ease of Use: A Challenge for the Industry](#)
[Connecting "Beyond the Spec"](#)

Ease of Use Building Blocks:

[Ease of Use Starts with Legacy Removal](#)
[Concept Platform Accelerates Legacy Removal at IDF](#)

Fall '99 IDF Articles:

[Intel gives a peek at future PC designs](#)
[When a PC becomes part of the furniture](#)
[Intel Developer Forum Keynote Speech, by Pat Gelsinger](#)
[Intel Announces New FlexATX Motherboard Form Factor for Low-Cost, Easy-To-Use PCs](#)

Other Related Links:

[Instantly Available PC \(IAPC\)](#)
[USB](#)
[Motherboards](#)
[Chipsets](#)
[Clarifying the Digital Display Interface Picture](#)
[Reducing Total Cost of Ownership with Intel® WfM and Microsoft ZAW Initiatives](#)

Designing Next-Generation Performance Desktop Systems

Joe Van De Water
Marketing Manager, Desktop Chipset
Platform Components Group
Intel Corporation

To take full advantage of the Internet, businesses and consumers require solutions that provide secured connectivity, scalability, accuracy, and responsiveness. The personal computer will be the 21st century's linchpin for meeting these demands. This fall, Intel will introduce a new architecture to meet these expanding requirements.

To address performance, all the system buses will be widened—SDRAM to RDRAM, AGP2X to AGP4X, the processor bus from 100 MHz to 133 MHz, ATA33 to ATA66, and I/O available bandwidth to memory from 133 MB/s to 266 MB/s. To meet the needs of manageability and security, we add self monitoring and alert across the LAN capabilities, while pumping up security with primitives built directly into the hardware. For ease of use, the PC will become instantly available with quick wakeup. It will also become more reliable with the elimination of legacy buses like ISA. This new platform will be built around the Accelerated Hub Architecture found in Intel's next-generation chipsets, providing all I/O with a direct connection to main memory.

Running RDRAM double pumped at 400 MHz, operating AGP4X graphics at quad pumped 66 MHz, and compressing/decompressing encrypted data over the Internet, among other chores, are incredible challenges for the board and system designer of the new platform. Tools and methodologies of the past simply won't work to give the accuracy required to build these next-generation systems.

At the upcoming Intel Developer Forum, we've dedicated the Desktop Track and Desktop System Technology Labs to providing you the information you'll need to build these next-generation platforms. We've lined up the experts within Intel, so we can unveil the details of this new architecture, describing how to build, test, and tune systems based on the upcoming chipset from Intel. Our format will be both class lectures and lab demonstrations to show and tell you how to bring it all together. Here's an outline of the two-day desktop track and three-day desktop lab track; more details on the session abstracts and registration information can be found at the [Intel Developer Forum](#) Web site.

Desktop Track: Designing Next-Generation Performance Desktop Systems

The desktop track begins with Richard Malinowski, engineering manager, Platform Components Group (PCG), unveiling the new architecture, giving an update on RDRAM, and providing an overview of the desktop sessions. Following, Mayne Mihasci, system architecture manager, OEM Platform Solutions Division (OPSD), will provide a one-year view of the form factor roadmap, including the impacts of future IA32, FlexATX, and trends in small form factors. Then Tony Shaberman, applications engineering manager, PCG, will detail the new features and capabilities of the I/O Controller Hub in Intel's new chipsets, including AC97, ATA66, LPC, system management, and Firmware Hub. The first day will end with Pete Mueller, applications engineering manager and Marcus Grindstaff, applications engineer, PCG, providing a detailed description of how to design and build platforms based on Intel's next-generation chipset and RDRAM.

The next day will begin with a session looking out at the technical platform challenges over the next two years. Scott Noble and Howard Heck, senior industrial designers in Intel® Architecture Labs, will cover topics from power supply architecture to new chassis discussions to the 2001 performance desktop platform. Afterwards, David Perchlik, system development engineer, OPSD, will introduce a roadmap for power supplies for performance and value processor desktop designs. Larry Ferra, products engineer of PCG, will then provide instruction on utilities Intel has developed for designing with Intel's graphics solutions, both discrete (Intel® 752 graphics accelerator) and integrated (Intel® 810 chipset). The desktop track will end with Frank Hady, staff system architect, Performance and Architecture Labs, demonstrating a tool to measure the concurrency and bandwidth capabilities of the new platforms and mapping these capabilities onto the latest software and benchmarks.

Desktop System Technology Lab Track

Measuring and simulating AGP4X and RDRAM signals will require new tools and methodologies. To introduce these tools and demonstrate how to use them on these next-generation platforms, we'll have three hands-on labs.

Validating Direct RDRAM and AGP4X with HP Logic Analysis System, will include an overview with demonstrations of techniques for measuring Direct RDRAM and AGP4X signals. Elizabeth Petropoulos, program manager at the Logic Analyzer Division of Hewlett-Packard, will demonstrate an RDRAM/AGP4X bus analysis tool that will aid in development of the upcoming high-performance desktop PCs.

Accurate device modeling and simulation is essential for designs with current and future generations of processors and chipsets. Development and Validation of IBIS Models, taught by Bob Ross, modeling engineer and Chair of EIA IBIS Open Forum at Mentor Graphics, will present through lab examples, applications and simulation results using IBIS features helpful in designing platforms based on Intel's upcoming chipsets.

New processors and chipsets coupled with new form factors will require different thermal solutions moving forward. In How Hot is Hot?, Jim Shipley, mechanical engineer in the Desktop Systems Development group of Intel's OEM Platform Solutions Division will provide interactive demonstrations and lecture on thermal test and evaluation methodologies.

Performance, manageability, security, and ease-of-use—all are requirements for the platform moving into the 21st century. The new platform brings many new challenges for the board and system designer, but even more opportunities in system design and flexibility. Register for the IDF desktop track to learn more about the next-generation platform architecture and how to develop scalable, optimized PC systems.

About the Author

Joe Van De Water is desktop chipset marketing manager for Intel Corporation's Platform Components Group. Joe has worked at Intel for the past 7 years, holding several positions within the Platform Components Group ranging from product planning to product marketing. Prior to Intel, Joe worked as a manufacturing engineer at Raychem. He holds a bachelor and master degree of science in Industrial Engineering from Stanford University.

For More Information

[Intel Developer Forum](#)

[Easy PC](#)

[RDRAM](#)

[Chipsets](#)

eHomes: Internet Connectivity Throughout the Home

Rochelle Keeler
Product Marketing Engineer
Intel Architecture Labs
Intel Corporation

More and more consumers are getting plugged in to the Internet at home. To enhance their Internet and computing experience, the challenge is to get them high-speed Internet access, a home network infrastructure that distributes the Internet throughout the home, and Internet appliances that can extend the usefulness of the PC throughout the home. The eHome is a reality, and developers need to understand the technology, consumer, and home environment in order to make decisions. The "Accelerating Broadband Internet Deployment to and Within the Home" track will address these issues at the Intel Developer Forum August 31 through September 2 in Palm Springs.

Consumers are demanding Internet connectivity in their homes. According to International Data Corporation (IDC), the number of PC-owning households going online is increasing. In their report "The Online Nation," IDC states that by mid-1998, 44 percent of households owned a PC, and 55 percent of those households were online. Twenty-five percent of all households were online versus 20 percent only six months before.

While consumers value their PC Internet experience, a PC located in the office or den will cause the family PC users to isolate themselves in these rooms. However, ethnographic research has shown that families want to spend time together in the home and are interested in products and services that meet this need.

Intel is exploring ways to enhance the Internet experience by enabling high-speed access to homes and unleashing the power of the PC anywhere in the home. The approach is threefold: working to eliminate the barriers to broadband deployment; providing the standards necessary to enable home networking technologies; and developing the concept devices and applications for Internet appliances in the home.

The Intel Developer Forum will have a track focused on computing in the home. The "Accelerating Broadband Internet Deployment to and Within the Home" sessions delve into the issues that conference attendees need to understand to help grow the home computing market. These issues include:

Broadband Deployment

Session One

Broadband Deployment Challenges and Opportunities

Tuesday, August 31, 10:30 a.m.-12:30 p.m.

This panel discussion includes participants from leading broadband hardware manufacturers, cable operators, telecommunications providers, and Internet Service Providers (ISPs). During this two-hour session, broadband expert Teri Lasley will encourage dialog about a number of broadband deployment issues:

- **Infrastructure**—What does the growth rate look like for the number of homes with cable or DSL (digital subscriber line) access? What resources are being allocated to stimulate this growth? What percentage of homes passed do you expect to reach by 1Q2000? What issues remain with deployment? Which technology offers a more aggressive cost reduction roadmap?
- **Retail and Standards**—How standard are the G.Lite and Data Over Cable System Interface Specification (DOCSIS) technical specifications? What is the consumer model for selling cable and DSL modems? What is the value proposition for the channel (retail, reseller, direct)? How will consumers know that they can get cable or DSL service at their home or office if they purchase modems?

- **Consumer Installation and PC Integration**—What is the consumer model for installing DSL or cable into the home? What service/support problems can be expected from cable and DSL and how can they be handled? How will the integration of modems into PCs help consumer installation?
- **Marketing and Content**—Which broadband technology will create more consumer pull? How do you get your marketing messages and benefits to consumers? What content is compelling to consumers with a broadband connection?

Session Three

Customer Premise Equipment (CPE) Controlled Cable Modems

Tuesday, August 31, 4:45-5:30 p.m.

Today's PC already contains the processing power, memory, and IP (Internet protocol) stack support needed to handle many cable modem functions. Therefore, many non-real-time cable modem functions can be implemented in software on the PC, resulting in a CPE Controlled Cable Modem. This presentation will provide the technical details and implementation strategy, and address cable operators' concerns.

Session Four

Implementing G.Lite and Discrete Multitone Technology (DMT) to the Home

Tuesday, August 31, 5:30-6:15 p.m.

G.Lite (G.992.2, or Universal ADSL) is poised to deliver Internet services "the last mile" over the copper telephone network. This presentation will describe DMT and how to implement this splitterless technology in the home.

Session Five

Communications Diagnostics

Tuesday, August 31, 6:15-6:45 p.m.

Historically, the communications subsystem has required a great deal of technical support due to both user and system issues. This presentation will focus on the causes of this support requirement and will introduce, describe, and demonstrate methods to reduce this support burden.

New Devices and Applications

Session Two

How to Build a PC-Connected Internet Appliance

Tuesday, August 31, 2:30-4:30 p.m.

Lightweight, portable Internet appliances that take advantage of the Internet connectivity and processing power of the PC will allow consumers to have Internet access anywhere in the home. This session will cover appliance design implications based on the consumer requirements for these products in the home. Using an Internet tablet case study, additional presentations will explore hardware design and software architecture tradeoffs to consider when creating this device. The concluding presentation will address the content and services that are key to the success of Internet appliances.

Home Networking Technologies

Session Six

Bridging the Gap Between HomeRF and HomePNA

Thursday, September 2, 2:15-3:15 p.m.

Standards-based home networking products are the key to unlocking the mass consumer market segment. This presentation will provide a brief overview of two home networking technology standards, HomeRF for wireless networks and HomePNA for phoneline-based networks. The bulk of the session will address bridging requirements between these two technologies and implementation challenges.

Session Seven**Consumer Gateways**

Thursday, September 2, 3:15-4:15 p.m.

Multiple PC homes and the need for home networks mean that consumers will soon face the need for a device that sits between the home network and the Internet. A consumer gateway enables multiple devices in the home to share a single connection for Internet access, and may also act as a firewall, provide Virtual Private Network (VPN) support, and enable IP telephony. This session examines the importance of consumer gateways, describes the implementation challenges, and offers suggestions on making a device that is easy for consumers to use.

The "Accelerating Broadband Internet Deployment to and Within the Home" track is a great way for a developer to get the latest information on the eHome, including broadband deployment, home networking technologies, and Internet appliances for the home.

More About the Author

Rochelle Keeler is a product marketing engineer for the Connected.Home initiative in the Intel Architecture Labs. Her group evangelizes Intel's vision for home computing, home networking, and high-speed Internet access. It drives the development of industry specifications for home network communications and the enabling technologies required to make the vision a reality.

Rochelle has been with Intel since 1998. She holds a Bachelor's degree in Anthropology from Pomona College and a Master of Business Administration from the University of Washington.

For More Information

For information on Intel's Connected.Home initiative, visit our [IAL Connected.Home Initiative](#) Web site. To register for this year's Intel Developer Forum conference, visit the [Intel Developer Forum](#) Web site.

Internet Driving Change in High Volume Servers

Justin Rattner
Intel Fellow and Director
Server Architecture Laboratory
Intel Corporation

Major developments in the fast-paced volume server industry are again the focus of the Server Track at the Fall '99 Intel Developer Forum. From round-the-clock Internet availability to the latest on VI architecture, attendees will be immersed in the newest technology and hottest trends in server architecture and design.

Our first day begins with the track keynote entitled **The Changing Face of the High Volume Server**. Rapid evolution is at hand as the inexorable force of the Internet mutates the high volume server into an array of new shapes and sizes, and the physical changes force a re-examination of everything we've come to assume about the classic SHV server. We'll take a fresh look at these changes and the impact they'll have on the volume server industry.

Leading off the Server Track courseware is a timely update on the **Portable Simulation Framework (PSF)**. Industry leaders, including Intel, IBM, LSI Logic, Unisys, and SES, have been quietly testing and evolving advanced system simulation concepts for this breakthrough effort. Three proof-of-concept studies, featuring the mixture of simulation models from different companies, establish that the PSF is now ready for broad industry engagement.

Day 1 continues the latest thinking on **Volume Servers in the 24x7 World**. Speakers from Intel will overview the changing industry priorities and highlight an effort to build a strong base of quantitative volume server availability information. Microsoft will cover the key changes in Windows* 2000 aimed at increasing server uptime in real-world applications, and IBM will tackle the challenging issue of a common framework for both online and offline diagnostics

Day 1 concludes with an ever-popular topic in the Server Track: system management. The session will start off with brief updates on the **Intelligent Platform Management Initiative (IPMI)** and the **Intelligent Platform Management Bus (IPMB)**. The session will then shift gears to introduce a new management interface, called **Metolious 2**, which brings ACPI-compliance to IPMI and other hardware abstraction layers.

Day 2 is all Server I/O. Things will get off to a rousing start with a multi-vendor discussion of **RAID Trends and Options for Entry and Mid-range Volume Servers**. Experts from Dell, DPT, Intel, and LSI Logic will offer their insights into what's required to stay competitive in this fast moving field. Audience members will get their chance to challenge the conventional wisdom in rapid-fire panel discussion to close out the session.

Another fast moving server I/O area is the **Next Generation Input/Output (NGIO)** specification and the 80-member strong **NGIO Forum**. Industry leaders will be on hand to present an update of NGIO Forum developments, including the just released NGIO Revision 1.0 Specification. NGIO specialists from across the industry will share the latest in NGIO product plans and demonstrate the first operational NGIO-compliant hardware.

Day 2's I/O theme continues with a course on **High Performance PCI Implementation** issues and answers. Focusing on 64-bit, 66-MHz PCI designs, the PCI practitioner will get a full dose of do's and don'ts to heed when designing new adapters and controllers for the high-speed I/O bus. A companion course will take a top-down view of **High-throughput Server I/O**. With volume server I/O hardware now capable of 1 GB/s of bandwidth, application developers will learn how to turn this raw power in benchmark-winning performance.

Server excitement continues unabated on Day 3. We'll hit the ground running with an industry power breakfast entitled **Enhancing Linux on IA for the Enterprise**. The impaneled movers and shakers of the Linux world, including Red Hat, Caldera, VA Linux Systems, and SuSE will tackle the tough issues of turning the upstart operating system into an enterprise-worthy platform. After opening position statements, audience members will be encouraged to get in the game with their most challenging Linux questions.

The afternoon of Day 3 promises to keep things as hot inside as it's sure to be outside when the Server Track tests the heat in the volume server segment known as **Server Appliances for Small Business**. Intel and Microsoft will present a fusion of requirements and key technologies for building very low cost, feature-rich servers for small and very small businesses. If you don't have a server appliance on your roadmap today, you'll surely be adding one after attending this session.

Day 3 keeps pulling things out of the technology furnace by getting up-close and personal with the newest member of Team SHV, the 8-way IA-based Pro-Fusion server. Intel experts will take a detailed look at what has this killer 8-way architecture burning up the benchmarks and offer timely insights on how to get the same results on your most important Internet applications.

The Fall IDF '99 Server Track concludes with an up-to-the-minute look at the latest advances in Virtual Interface (VI) Architecture by **Introducing the VI Developer Forum**. With VI architecture an integral part of both the NGIO and Future I/O (FIO) specifications, hardware and software developers will both want to hear how industry has come together to make the VI Architecture easier to implement and easier to use for both I/O and clusters.

From advanced server simulation technology to tomorrow's low-cost server appliances, if it's happening in servers, it's teed up and waiting for you at the Fall '99 IDF Server Track in Palm Springs, California, August 31 through September 2, 1999. We'll see you there!

About the Author

Justin Rattner is an Intel Fellow and Director of Intel's Server Architecture Laboratory. His current R&D activities focus on future generation IA-32 and IA-64 server technologies and standards. In December of 1996 Justin was featured as Person of the Week by ABC World News for his visionary work on the Department of Energy ASCI Red System, the first computer on earth to achieve sustained performance of one trillion scientific calculations per second. In 1989, he was named "Scientist of the Year" by R&D Magazine for his leadership in parallel and distributed computer architecture. Justin was recently honored as one of the Computing 200, the Computer Museum's project to recognize the 200 leaders having the greatest impact on the U.S. computer industry today.

Justin joined Intel in 1973. He was named its first Principal Engineer in 1979 and its fourth corporate Fellow in 1988. Prior to joining Intel, Justin held positions with Hewlett-Packard Company and Xerox Corporation. He received the BS and MS degrees from Cornell University in Electrical Engineering and Computer Science in 1970 and 1972, respectively.

Applied Computing Firmly Embedded at IDF

By Will Schreiber
Senior Systems Engineer
Computing Enhancement Architecture Lab
Intel Corporation

"Applied computing" refers to the high performance, dedicated and connected portion of the embedded market segment. Intel products are helping fuel significant growth in applications and capability within this space. Areas of expansion include voice/data communications platforms, retail and financial transaction terminals, industrial PCs, and wireless and remote-access products.

At the upcoming Fall '99 Intel Developer Forum a technical track dedicated to applied computing provides a detailed exploration of how Intel® StrongARM® processors and Pentium® II processor-based platform technologies can provide a fast time-to-market foundation for devices that enhance networked information systems.

Intel® Architecture

The vast developer knowledge base and the super-abundance of development tools and Intel® Architecture (IA) application software can dramatically accelerate time-to-market for applied computing devices.

In the applied computing track, one course provides developers with a new low cost option for applied computing applications. The course is an in-depth discussion of how to fast-boot IA applications without the added cost and overhead of a BIOS by using the new Intel® System Software Library (SSL).

Another course details the power management advantages of Intel Architecture platforms running the Windows® CE operating system. Using the power management features built into Intel® chipsets, Intel worked directly with Microsoft to extend and enhance the power management capabilities of Win CE. The use of this advanced power management software and its interface to your application is the focus of the course.

While the Pentium II processor can run code generated for earlier Intel Architecture processors without modification, new compiler enhancements take full advantage of the processor's superscalar architecture, and increase your application performance. Another applied computing course teaches application developers how to get maximum performance from their application running on Pentium II processor platforms, by using the GNU compilers from Cygnus Solutions*. Course topics include increasing superscalar efficiency with better code alignment, boosting floating-point performance, and how to use the new Pentium II processor instructions to speed execution.

For Industrial PC designers, the applied computing track also includes a course about Intel Architecture hardware design using the *Embedded PC 2000 Design Guidelines*.

Intel® StrongARM® Technology

Because of its high performance and extremely low power consumption, the Intel StrongARM processor family is the low-power processor of choice for handheld or connected applied computing devices. The StrongARM processor combines the highest performance at the lowest power consumption of any processor available.

An applied computing course you don't want to miss is the architectural overview of the next-generation Intel StrongARM processor. The course explains the design of the Intel® Super-Pipelined RISC Microarchitecture of the next-generation Intel StrongARM processor. This design extends processor performance to 1700 to 4500 MIPS per Watt while consuming lower power per MIP. With the optimal power-performance-cost ratio for both handheld and wired applications, this next-generation Intel StrongARM processor enables applications that were once designed for desktop PCs—to run on a handheld device.

The applied computing track also includes a “cookbook” course on how to port the Linux* O/S to StrongARM processor-based platforms. Included in the course are ways to modify the kernel source code to accommodate different physical memory configurations and device drivers. The course also covers how to put together the cross-development tools needed to build the Linux kernel for the SA-1100 target, as well as a list of Linux resources.

Web-based management of wireless devices is the focus of another course in the applied computing tract. The course shows how a network administrator can quickly and remotely configure an appliance. The course presents an architectural overview and case study using the Web to configure and manage network devices with a simple and friendly Linux browser. It features RapidControl* products from Rapid Logic Inc.*.

About the Author

Will Schreiber is a Senior Systems Engineer in Intel’s Computing Enhancement Architecture Lab. The lab is one of several labs at Intel focused on the future of computing and networking technologies.

For More Information

- For information about the Fall'99 IDF Conference, including technical track information and registration details, visit the [Intel Developer Forum](#) Web site.
- For information on applied computing platform solutions, visit the [Intel Platform Solutions](#) Web site.
- See the article on [Intel Next-Generation StrongARM Technology](#) by Allen Hyman in the May issue of *Platform Solutions News*.

Note: ARM and StrongARM are trademarks of Advanced RISC Machines, Ltd. Third-party brands and names are the property of their respective owners.

Inside Looking In

Of Technology and Fruit

By Tim Mostad
Senior Technical
Marketing Manager
Intel Corporation

Puzzled about how to develop and then profit from new technology? Our resident philosopher and technical enabler, gentleman farmer Tim Mostad, can show you where the road to success leads. Take an enlightening trip from the barn, through the orchard, and back again for a slice of warm, homegrown insight.

When I was young, my mother would rhetorically ask, "What? Were you born in a barn?" She probably was just trying to tell me to keep the door closed. She grew up on a farm but I knew nothing of farms and farming, so I never appreciated the underlying contradiction. Now that I own a farm with a barn and actually try to grow things, the question has new meaning. Moreover, I realize that it now applies to me and to my department's mission at Intel.

Developing and then profiting from new technology is quite a bit like farming. In both processes, you start with virtually nothing but a patch of ground you call your own. Whether this is literally dirt or just an area of technology where you've staked out expertise, everything you accomplish relies on creating a clear space. Intel's farmers live in our Intel® Architecture Labs, and we work closely with them to deliver technologies we believe can feed the industry with clear benefits.

I manage a small group of technical marketing engineers whose job it is to make our customers' lives easier as they attempt to follow Intel® platform technology roadmaps. We harvest the key learnings from the technology architects, process them, then ship them in the form of design collateral to our customers. Similarly, PC product developers take the new platform technology we produce, package it in crates that we commonly call PCs, and deliver consumer-ready products to market. The epitome of our success as TMEs would mean the product development process would be no more difficult than one more crank of the gristmill.

However, unlike flour, the technology we market from the Intel Architecture Labs gristmill has a value that's often hard to quantify. In sharp contrast to farm produce, which is most desirable at the peak of ripeness, new technologies need to be sold when they're still green. These technologies have to precede observed problems by a sufficient amount of time to be integrated, tested, and made ready for the end user. It's often difficult to get anyone to design with a new technology for which there's no obvious market or requirement.

On my real-life farm, for example, I could grow a "squidgy fruit" but who'd buy it? Aside from the fact that I just made up the name and no one could possibly know what it is, I'd have a huge job ahead of me if I wanted to make money growing a product no one recognized.

I'd need to create a squidgy fruit market. This would involve a tremendous amount of consumer education, including comparisons to competing produce such as raspberries, bananas, and avocados. I'd also need to publish delicious, easy-to-fix recipes using this succulent, edible, pulpy mass. Knowing of an alternate use, such as replacing coffee filters with dried squidgy fruit skins, certainly wouldn't hurt my potential for acceptance either. On it would go as I tried to establish a clear identity for my produce. In this scenario, it's not so important that the fruit of my labor is perfect; it's more important that it's clearly wanted.

Such is the world of electronics technology as well. The PC platform landscape changes constantly and rapidly. Companies like Intel, which have a huge investment in pushing the technology envelope, need groups like mine to help keep our technology consumers onboard. We spend the majority of our time differentiating our technology for them, like the guy showing the Popeil pocket squidgy fruit juicer at the local county fair. If our demos at tradeshow and conferences aren't convincing, then the market won't buy our technology, either.

The difficulty with creating new markets is that it's incredibly tough to establish and maintain the ideal position, where you're the only supplier. No sooner do you perfect your squidgy fruitcake recipe—and just in time for the Christmas fruitcake rush—than you see an ad on TV from Time-Life Books for the “Best of Squidgy Fruit Cake Recipes from the 60s, 70s, and 80s.” In reality, nobody wants to be dependent on you for the sole supply of his or her essential nourishment, and you really don't want that role, either. You just want your customers to like your formula best.

People inside Intel don't always understand this too well. As a natural consequence of enabling designers to rapidly assimilate technology, we sometimes have to attract actual or potential competitors to our enabling events, such as plugfests. On one occasion my group allowed a direct competitor to sponsor a very nice after-hours social event. We were accused, in essence, of being born in a barn. We had left the door open, inviting a competitor to step in. To hear the complaints you'd have thought we'd given away Grandma's secret recipe for squidgy fruit compote. Talk about sour grapes!

In fact, we all know how nimble and talented our competitors are, and we know that any advantage we have won't last. That's why we continue to push the bounds of what's possible, looking for that secret ingredient that makes our products better. The good news is that the ultimate winner is the consumer, not of some fictitious food, but of real PC products. The computing power and features that the PC industry delivers for the money is ever more impressive and no end is in sight.

Farmers routinely leave the barn door open because they have a real sense of security. When we work to create new markets for technology, we have to have the confidence that we can create the world's best products based on the world's best technology to the benefit of our bottom line and our shareholders. Competition cannot be a reason to shut out the world. Instead, it needs to be a reason to redouble our efforts, to focus. It fuels our sense of urgency and keeps us on our toes. Not only that, it's inevitable. So, Mom, I hope we'll leave the barn *and* technology doors open and show some good old county hospitality by inviting our neighbors in for a big slice of squidgy fruit pie.

About the Author

Tim Mostad says, "the majority of my 18 years at Intel have been spent in the pursuit of technical marketing nirvana." He is responsible for demos, white papers, plugfests, and technical training to support the adoption of new desktop technologies.

Top Stories

Direct DRAM Update: Test Results Available on the Web

By Peter Mueller
Platform Architecture Manager
Platform Components Division
Intel Corporation

Because there's no margin for error with Direct RDRAM devices, Intel established a third-party testing program to ensure that RDRAM components and RIMM modules meet the Direct RDRAM Specifications. Compliance testing results for participating vendors are in—and out on Intel's Chipset Web site.

For several years, Intel and Rambus, Inc. have worked together to develop a new memory technology for PC platforms known as Direct RDRAM. PCs equipped with Direct RDRAM have significantly more memory bandwidth compared with the current generation PC100 SDRAM-based systems. Moreover, one of the principal advantages of Direct RDRAM is the ability to let applications scale in performance with faster processors, graphics, and I/O for years to come.

Direct RDRAM components, RIMM (Rambus Inline Memory Module) modules, DRCG (Direct Rambus Clock Generator) clock components, and RIMM connectors have passed validation testing to production-level specifications. This step sets the stage for the industry to launch a next-generation PC platform with the performance and scalability benefits of Direct RDRAM. We expect Intel's next-generation chipset, scheduled for introduction in September, to be the first available to support Direct RDRAM.

Validation Benefits

Direct RDRAM specifications are considerably more exacting than SDRAM. Because there's no margin for error with Direct RDRAM devices, Intel has made a considerable investment to establish a third-party testing program to ensure that RDRAM components and RIMM modules meet the Direct RDRAM Specifications.

What's more, Intel is making its test results openly available on the Web. With the publication of test results, Intel is providing an incentive for manufacturers of quality Direct RDRAM products to participate in the program, as well as providing OEMs with important validation data on specific products. Compliance with the specification is a huge step toward the widespread deployment of Direct RDRAM.

The Results Are on the Web

Direct RDRAM compliance testing results for participating vendors are now posted on the Direct RDRAM page of Intel's Chipset Web site. When you visit the site, you'll understand why, at Intel, we're extremely pleased with the results.

Here's what you'll find posted:

- **Direct RDRAM Component Testing Summary, Revision 4.0**—Intel has been working with RDRAM vendors and external testing companies to measure the performance of each RDRAM vendor's 400 MHz RDRAM components against the requirements of the Direct RDRAM Specifications. The Web site features a summary of the validation results.
- **RIMM Module Validation Results**—Intel has also worked in cooperation with RIMM vendors and independent testing organizations to measure each vendor's 400 MHz RIMM performance against the requirements of the Direct RDRAM Specifications. The Web site provides a downloadable file with a summary of the validation results.

- **RIMM Module Reference Designs**—You can download Gerber files for the implementation of RIMM module designs. Continuity RIMM (C-RIMM) Module Reference Designs and a Production Heat Spreader Reference Design are also available for download.
- **Direct DRAM is On-Track**—Intel has demonstrated this new Direct RDRAM memory technology at work in a PC platform. By successfully showing the interoperation of Direct RDRAM technology ingredients including Direct RDRAMs, RIMM modules, RIMM connectors, clock chips, and motherboard design and layout, Intel has confirmed that the advent of Direct RDRAM-based systems is a reality.

The Direct RDRAM program is currently on-track for the September announcement of chipsets that will bring Direct RDRAM to the desktop, workstation, and server market segments. These chipsets will enable the processor, graphics, and I/O subsystems to take full advantage of the higher Direct RDRAM performance.

We encourage developers to visit the Web site and become familiar with the Direct RDRAM suppliers who meet the compliance specifications and who are now ready for production. Developers should obtain and sign the reciprocal, royalty-free patent license for Adopters of the Direct Rambus RIMM Specifications. The agreement is available on Intel's RDRAM Web page.

The Direct RDRAM specifications are available for download directly from the Rambus Inc., Web site.

With this head start on understanding Intel's next-generation chipset, developers will find all the pieces in place for volume production of a new generation of scalable, high-performance PCs.

For More Information

Visit the [Intel® Chipsets RDRAM](#) Web page for additional information on Direct RDRAM Component Validation Results, RIMM Module Validation Results, RIMM Module Reference Designs, Direct DRAM Specifications and links to application notes. You can also register to receive Direct RDRAM updates.

Visit the [Rambus, Inc.](#) Web site to download the Direct RDRAM specifications.

About the Author

Peter Mueller is a platform architecture manager in Intel's Platform Components Division, responsible for technical enabling and launch of new initiatives, including Direct RDRAM technology. Pete has been working at Intel since 1988 and has an MSEE from UC Davis.

NGIO Set to Offer Next-Generation I/O for the Internet

Mitch Shults
Director, NGIO Initiative Marketing
Enterprise Server Group
Intel Corporation

As customers begin to learn about NGIO, we can expect demand for products to kick in rapidly. This article reviews the benefits of NGIO, and talks about upcoming events where developers can get more detailed information.

NGIO is designed to overcome the limits of current bus-oriented I/O architecture. Instead of shared-bus, memory mapped I/O with a single fault domain (Figure 1, left), the NGIO architecture (Figure 1, right) offers a channel-oriented fabric of fast, inexpensive links that decouple I/O devices from the CPU.

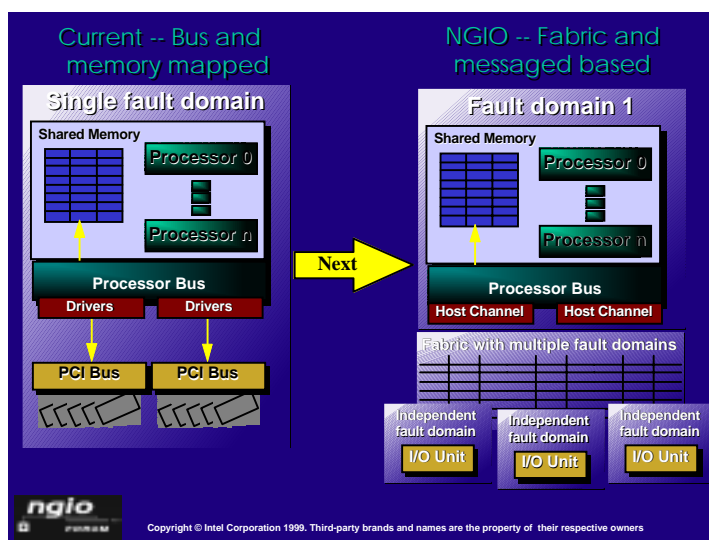


Figure 1: Shared-bus versus Channel-oriented Fabric

NGIO addresses the critical server bottleneck between server memory and I/O controllers (Figure 2) to provide the reliability, performance, scalability, and flexibility that today's I/O-centric Internet servers (and other servers) demand.




Figure 2: Server Bottleneck

Expect Customer Demand for NGIO

The old song asks, "How're you gonna keep 'em down on the farm after they've seen Paree?" For NGIO, the questions developers need to consider is how they're going to satisfy their customers with current-technology designs once businesses have gotten a glimpse of what NGIO can do for them. We anticipate significant demand for NGIO-aware products starting in the Year 2000 time-frame, as Information Technology (IT) organizations become aware of what NGIO can provide. Customers will demand NGIO solutions because NGIO enables higher levels of:

- **Reliability**—NGIO's point-to-point fabric connection allows natural redundancy and makes it easy to detect and replace faults. Because I/O is no longer mapped to system memory, device conflicts can't corrupt other areas of memory and crash the system.
- **Scalability**—NGIO scales easily with additional attachment points. It's designed to support high-density rackmounting.
- **Performance**—Each link runs at 2.5 gigabits per second, and—more importantly—links can be aggregated to multi-terabit throughput.
- **Flexibility**—NGIO can be used outside the server complex by adding extra switches, overcoming the slot limitations often encountered in today's systems. NGIO can also be used for server-to-server connectivity to provide further flexibility.
- **Protocol Flexibility**—NGIO can accommodate all higher level software protocols (IPV4, IPV6, SNA, IPX/SPX, etc.). You (and your customers) aren't locked into one protocol, and can take advantage of advances in protocols without having to make expensive and time-consuming hardware changes. This is particularly important in the Internet arena, where protocols, applications and technologies are evolving rapidly.
- **Cost-Effectiveness**—Using just four pins (compared to 100 pins for a PCI connector), NGIO has been designed as an efficient solution that's affordable for entry-level servers while scalable to meet high-end demands.

Moreover, NGIO will be attractive to customers because of its timeliness. NGIO has been under development for over two years, and work on prototype boards is well underway. Customers need solutions today, and NGIO is the fastest way to provide those solutions.

Be Ready to Meet the Demand

NGIO is the I/O solution for future product designs, and now is the time to get involved in designing NGIO products, if you're not already. I urge you to:

1. **Get the specs**—As soon as they're publicly available, the NGIO specifications will be posted at the [NGIO Industry Forum's](#) Web site:
2. **Join the Forum**—Join the NGIO Industry Forum so you can influence the development of NGIO and work hand in hand with the companies driving the [NGIO](#) specifications:
3. **Attend upcoming events**—Intel representatives will offer a session on NGIO at [IDF](#) in August.
4. **The NGIO Industry Forum is sponsoring a Developers' Conference** September 28 - 30 in Newport Beach, Calif. To get more information and to register, see the [NGIO Forum](#) Web site.

About the Author

Mitch Shults is director of NGIO Initiative Marketing. Prior to his involvement in NGIO, Mr. Shults was the director of server platform marketing for Intel's Enterprise Server Group, responsible for coordinating the marketing for server industry technology development efforts.

For More Information

[Momentum Grows Behind Next-Generation I/O](#)
[Aberdeen Group white paper](#)
[NGIO Industry Forum](#)

Five Steps to Scalable Linux Applications

Tim Witham
Principal Engineer and Linux Engineering Manager
Intel Corporation

For Linux to become mainstream in the IT environment, it must be usable across an application mix that ranges from single-processor workstations to 8-way symmetric multiprocessing (SMP) systems. This article gives five simple rules that take developers a long way toward achieving scalable Linux applications.

Linux, already the leading operating system on the Internet, is increasing its market share in the desktop, workstation, and low-end server arenas and is rapidly entering the mainstream of the data processing world. That means it's moving onto systems with more processors and larger memories, often more than 4 Gigabytes of main memory.

Naturally, this puts a greater importance on scalability. For Linux to become mainstream in the IT environment, it must be usable across an application mix that ranges from single-processor workstations to 8-way symmetric multiprocessing (SMP) systems. While a single application mix might not run on these very different configurations, systems at either end of this spectrum will need to run Linux—and to run it well—if Linux is to be embraced by the data center. If Linux and its applications cannot achieve this scalability, there's a very real chance that it will be relegated to serving niche or appliance applications.

Fortunately, many scalability problems are easy to avoid. In this article, I'll give you five simple rules that will take you a long way toward achieving scalable Linux applications. But first, let's take a step back and consider one of the fundamental notions of scalability: Amdahl's Law.

Amdahl's Law of Speed-up

Gene Amdahl, the chief architect of the IBM System/360* computer and founder of Amdahl Corporation, promulgated Amdahl's Law in 1967. Amdahl's Law describes the amount of speed-up that can be achieved from a given improvement to a multiprocessor program by noting that the potential speed-up is limited by the fraction of the code that must be performed sequentially. He quantified this as:

$$Speedup = \frac{1}{((1 - FractionEnhanced) + FractionEnhanced/SpeedupEnhanced)}$$

Figure 1. Amdahl's Law

For example, if 20 percent of a program must be performed sequentially, then the maximum speed-up is fivefold. The larger the proportion of a program that must be performed sequentially, the more limited the potential speed-up.

For a multiprocessor system, code that is locked or is blocked because of data misalignment cannot be executed in parallel. It thus limits the scalability of the application. To see how quickly this exclusion can prevent an application from scaling, let's look at the following table.

	Processor Count		
% blocked	2	4	8
0.01%	2.00	4.00	7.99
0.10%	2.00	3.99	7.94
1%	1.98	3.88	7.48
10%	1.82	3.08	4.71
50%	1.33	1.60	1.78
90%	1.05	1.08	1.10

Table 1. Parallel Execution Failure Table

As you can see, if just 10 percent of an application's total execution is blocked, scalability is significantly affected by the time four processors are brought into play. While this might seem like a lot of time given the size of the critical regions, it's important to note that this applies to execution time, not instruction counts. Because of caches and non-uniform memory accesses, the execution of critical lock sections that involve exclusive use of code or data can take much longer to execute than non-locking regions.

Five Simple Rules

On a single-processor system, these blocking issues would not cause problems. Often, however, coding practices that work fine with a single stream of an application can quickly create problems when multiple tasks are executing simultaneously. Through careful coding practices and system setup, you can avoid most of them. Here are some simple rules to give you a start:

1. Lock data, not code.
2. Never hold more than one spin lock at a time.
3. Keep data to be used with locks close to the locks.
4. Scatter locks.
5. Avoid the need to lock.

While there are many different ways of achieving good code that scales well, these five practices are easy to follow. In addition, they don't affect the running of the code on a single-processor machine, so they don't force you to maintain different versions of your code or follow different practices than for single-processor and multiprocessor implementations.

Lock Data, Not Code

Locking data instead of code gives you the ability to multithread hot routines. This will be a natural result of the individual request's having only the data that it is using locked. The only blocking functions that will occur are during the very short time that the process is determining the next record to work on, or when two or more processes actually want to work on the same element in a data structure.

One Spin Lock at a Time

Don't hold more than one spin lock at a time. This has two benefits:

- It makes it much harder to get into a deadlock situation with the very tight spin lock loops.
- By not holding more than one "fast" lock at a time, you eliminate the chances of some other part of the system delaying the release of the second lock and thereby increasing the hold time of the first lock to an unacceptable level.

Keep Data Close to its Spin Lock

Hot data locks are often passed back and forth between processors. This cache-to-cache transfer time is orders of magnitude longer than the time it would take the code to execute if it could reference data already in the processor's L1 or L2 cache. If the data the lock addresses would fit into the same cache line as the spin lock, you wouldn't have to do another cache-to-cache fetch. Instead, you could execute the critical code on data already in the processor cache. This could produce a much smaller critical region. This might not be noticeable on a single-processor system, but can be significant on a multiprocessor system since it means the hot data will still be in the single processor's cache.

Scatter Locks

Any implementation of "fast" lock should require that all lock data primitives start on a cache line boundary. This ensures that two hot records that are close in memory do not inadvertently result in a single locking data structure. If you require your data alignment to have all "fast" locks, begin on a cache line boundary. Then, by creating all data structures that have "fast" locks with the fast lock as the first element, you can ensure that you are not creating extraneous cache-to-cache transfers with corresponding increases in the critical regions.

Avoid the Need to Lock

If a producer/consumer relationship can be created with one element only writing and one element only reading from the queues, then the need for hot locks can be removed, with the resulting decrease in the critical regions

Craftsmanship, Not Magic

To recap, achieving scalable Linux applications is relatively easy if you follow some simple programming practices. Make sure not to hold application fast locks during OS calls, since this can make the critical region very large. The percentage of code that's not multi-threaded can be additive, which quickly mounts up. The percentage that isn't multi-threaded is calculated by runtime and not by instruction count.

About the Author

Tim Witham is an engineering manager in IPL. Tim has a broad range of experience with high performance commercial systems. This includes both benchmarks and real production "largest open system" including the creation of very large mission critical open systems. In the past he has held many performance records, including the first open system over 100 TPS and has over 20 years of experience with high-end commercial computer systems.

For More Information

At the upcoming Intel Developer Forum conference, I'll be offering a workshop on designing scalable Linux applications, where I'll discuss many of these topics further. For information on registering, visit the [Intel Developer Forum](#) Web site.

Increase Data Protection with the Hardware-Based Intel® Random Number Generator

By Clark Haass
Senior Product Marketing Engineer
Platform Security Division
Intel Corporation

Building trusted Internet applications is vital to the virtual enterprise. Keys used in cryptography and security protocols are only as strong as the random number sources used to create them. With ubiquitous hardware-based RNG, Intel provides software developers a readily accessible source of random numbers on the PC platform.

Random numbers are fundamental to good security. They serve as basic building blocks for the keys used in cryptography, digital signing and security protocols that are vital for trusted Internet applications. The industry currently uses pseudo random number generator algorithms* (PRNGs), for random number inputs. While good PRNGs employ cryptographically strong algorithms, they rely on seeding from a source of randomness on the client or server. The problem is that such commonly used "random" sources as keystrokes, mouse clicks, and hard drive timing data are not truly random and non-deterministic. This opens a security hole that can be exploited. Only hardware-based random number generation is capable of being truly non-deterministic or unpredictable.

Intel is providing the first mass-produced hardware-based random number generator (RNG), implemented in the Intel® 810 chipset. For the first time, software developers have a proven source of randomness on the platform they can use to strengthen the security of cryptographic functions.

Get off the Entropy Treadmill

To be cryptographically strong, PRNGs must be seeded with random and non-deterministic numbers. Even the best algorithmic sources for random numbers can be shown to be non-deterministic, prompting computer pioneer John Von Neumann's rather provocative remark, "Anyone who considers arithmetical methods of producing random numbers is, of course, in a state of sin."

That is strong language, but the fact remains that the security of data encryption, digital signing, and Internet security protocols, such as secure sockets layer (SSL), is built on the foundation of randomness, or entropy. Software alone does not provide true randomness. The search for natural sources of randomness puts developers on an "entropy treadmill," tapping into seemingly new random events as the frequency of keystrokes or OS system ID signals. It can be statistically demonstrated that such sources are not truly random.

The Intel® RNG enables developers to get off the entropy treadmill. Even more importantly, it provides a low-cost and ubiquitous source of high-quality, random numbers that can be used to build cryptographically strong keys for applications including Web browsers in clients and servers, virtual private networks, e-mail, e-commerce, Internet firewalls, and certificate authority applications. Other applications benefit from random numbers, including imaging, lotteries, and Monte Carlo simulations.

How Hardware RNG Works

The hardware-based Intel RNG is available in chipsets, beginning with the Intel 810 chipset. It is implemented in a new chipset component, the Intel® 82802 Firmware Hub Device (FWH). This component stores and manages system and video BIOS in flash memory. It also serves as the platform for security features, beginning with the Intel RNG, which uses thermal noise from a resistor (Johnson Noise) to generate a truly random, non-deterministic stream of bits.

Applications typically access the hardware-based Intel RNG through security middleware, such as RSA BSAFE* Crypto-C and Crypto-J, CDSA (Common Data Security Architecture), or Microsoft Crypto API. The middleware communicates with the Intel® Security Driver, which accesses the RNG in the chipset. The Intel Security Driver consists of two layers:

- An operating system-independent layer designed to insulate layers above it from changes in hardware;
- An operating system-dependent driver layer that accesses the hardware primitive.

The Intel Security Driver will be delivered through independent software vendors and PC OEMs.

Validation

Intel has validated the hardware-based Intel RNG throughout the design and manufacturing process. These steps include pre-design validation, export review with the U.S. Department of Commerce and European countries, post-design validation with Cryptography Research (CRI), and the Federal Information Processing (FIPS) Level-3 test for statistical randomness (FIPS 140-1). The Intel RNG is not subject to export restrictions, based on current regulations.

Summary

Building trusted Internet applications is vital for the continuing deployment of e-commerce, Virtual Private Networks, and other infrastructure for the virtual enterprise. Keys used in cryptography, digital signing, and security protocols are only as strong as the random number sources used to create them. Seeding PRNGs with numbers from truly random and non-deterministic sources can enhance the security of keys and the applications that rely on them. By delivering the first ubiquitous hardware-based RNG, Intel provides software developers with a readily accessible source of random numbers on the PC platform to improve the very foundation of virtually all Internet-based security applications.

About the Author

Clark Haass is senior lead PME in Intel's Platform Security Division. He and his team are focused on delivering security capabilities such as the Intel RNG to IA platforms as part of the Trusted Connected PC initiative. Clark brings over 7 years of product development and marketing experience and holds a BSEE, MSEE, and MBA. Also, Clark has been published on topics ranging from medical image processing interfaces to performance evaluation of print servers and he's an avid decorative balloon collector.

For More Information

Visit Intel's new RNG technology Web site for complete details, including a technical overview of RNG and its impact on security, technical papers, and documentation on the [Intel 810 chipset](#).

* The term pseudo random number generator is a term used by B. Schneier in his treatise on cryptography, Applied Cryptography.

Designing for High Server Availability

Sha Gleason
Server Availability Marketing Manager
Enterprise Server Group
Intel Corporation

Every piece of the solution stack has a role to play in achieving highly available servers, and every vendor needs to design for availability. This article takes a quick look at the issue of availability, including some of the work Intel's doing to ensure highly available Intel-based server solutions.

In today's nonstop world of electronic business, unplanned downtime can mean lost revenue, lost customers and interruptions in critical business processes. Not surprisingly, as the price and price/performance of Intel-based servers move them further into the data center and onto the Internet, the demand for high availability is increasing.

Every piece of the solution stack has a role to play in achieving highly available servers, and every vendor needs to design for availability. This article takes a quick look at the issue of availability, including some of the work Intel's doing to ensure highly available Intel-based server solutions. We're also offering a course at IDF in August on designing for server availability, and I'll tell you a bit about that as well.

Reliability + Serviceability

As Figure 1 shows, server availability is a combination of reliability (keeping systems up and running) and serviceability (getting them back online quickly when service is interrupted). Availability features range from reliable hardware and software to repair and upgrade technologies to best practices for deploying and managing servers.

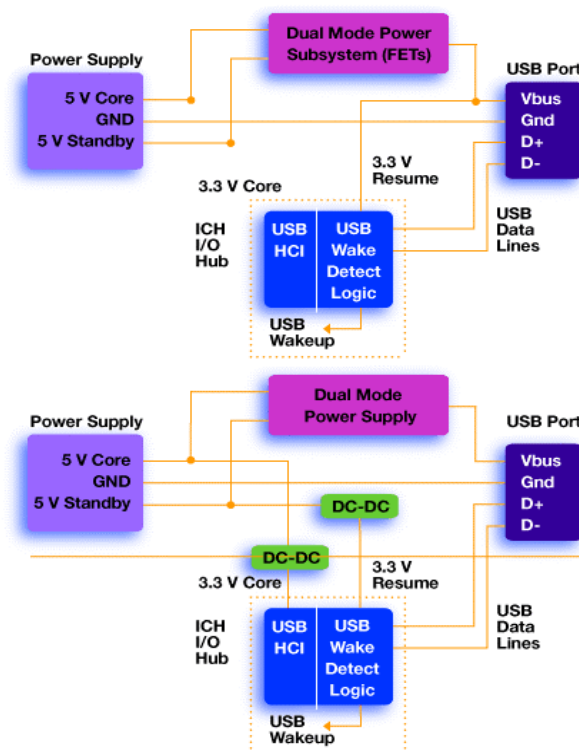


Figure 1: Server Availability

A recent study by the Gartner Group found that only 20 percent of unplanned server downtime results from technology (operating system, hardware, and facilities). The remaining sources are split fairly evenly between application failures and operator errors. Gartner further reported that much unplanned downtime can be “managed away” or avoided via improved training, careful IT processes, and automation.

The definition of availability depends on how a server is used. The goal is to meet the customer's requirements for the level of service needed without incurring unnecessary costs. Front-end servers generally achieve high service availability by failing over from one server to another if one server is down. In effect, they treat servers as field replaceable units, to keep services available to users. Back-end servers, with higher demands for single server availability, incorporate more costly highly reliable designs, greater fault resilience, and hot-pluggable subsystems. Fully fault tolerant systems (uptime of 99.999 percent or higher) incorporate broad-scale redundancies in their designs. While the market for fully fault tolerant computers is small and shrinking, the demand for server installation that provide 99.9-99.99 percent service availability is experiencing significant growth. Operating system clustering and other forms of clustering are increasingly providing this level of service availability.

What Does It Mean for You?

Availability isn't a function of just one aspect of server hardware design. Complete solutions depend on every level of the solution stack, which means that all of us have a stake in designing for high availability.

For comprehensive solutions, all levels of the stack must fit together smoothly and seamlessly, with each level providing the capabilities needed by higher and lower layers. That means it's not enough to ensure that our own “piece” of the solution stack meets availability requirements. To deliver robust, manageable solutions, we have to focus not only on high availability but also on the interfaces between our own availability features and the needs of the surrounding layers of the stack.

Customers don't want to have to manage each piece of a server solution separately, whether the machine is clustered or unclustered. They want to manage each server as a unit, mixing and matching components as needed. This places a high premium on standards and emphasizes the importance of developing and adhering to compatible ways of handling tasks such as diagnostics and monitoring.

Other implications:

- Platform and peripherals—Provide the appropriate amount of redundancy, along with well-designed monitoring and diagnostics tools.
- Operating system and software—Build in fault-resilience. Create operating systems, applications, utilities, and middleware that support software clustering, plus tools that enable IT departments to manage clusters as easily as if they were single machines.
- Service providers—Promote best deployment practices and policies to reduce avoidable problems caused by training and process deficiencies.

What Intel Is Doing

Without detailed real-world data about the root causes of system failure, developers must rely on theoretical failure prediction models that are notoriously conservative. To provide real-world information, Intel's Server Architecture Lab has created a RAS Measurement Program that involves instrumenting our corporate servers and analyzing failure data. This information will provide a base for developing strategies to further boost server availability.

Diagnostics are a key ingredient in the recipe for high availability. Without comprehensive online and pre-boot diagnostics, it's hard for developers to predict how their products will perform and for IT personnel to tell why a system has crashed, let alone head off problems before they occur. Intel is developing a Common Diagnostics Framework (CDF), a set of driver diagnostics and instrumentation that can be used for predictive analysis and fault/failure diagnosis. In keeping with our commitment to open industry standards, our engineers are working with the Desktop Management Task Force to make CDF available to developers who wish to use it.

Learn More at IDF

You can learn more about ensuring highly available server solutions by attending the Server Availability class at IDF in August. Along with information about CDF and other tools and techniques to help design highly available server solutions, Intel will preview its future server availability direction and Microsoft will discuss reliability improvements in Windows 2000* and Windows NT*.

With our combined efforts, we can all enjoy the expanding demand for highly available Intel-based servers—and IT organizations and end users can enjoy powerful, cost-effective and highly available servers for the full range of their performance needs.

About the Author

Sha Gleason is marketing manager for Intel's Server Availability Program. During her 13 years at Intel, she has migrated from processor component design engineer and design manager, through technical marketing and into marketing programs. Sha holds an M.S. in Electrical Engineering and a B.S. in Applied Mathematics. When she's not at Intel, you can find her at her loom or harp.

Digital Visual Interface Comes into Focus

Carol Jacobson
Initiative Technical Director
Intel Corporation

The Digital Visual Interface (DVI) specification created by the Digital Display Working Group is designed to stimulate the widespread adoption of digital displays for high-performance desktop and mobile PCs. Find out how DVI offers a new way to provide display products at lower cost.

The PC industry is at long last "going digital," all the way from the graphics subsystem to the display.

Until today, the lack of a standardized digital interface, coupled with the relatively high price of digital displays, has hindered the adoption of all-digital PC display technology. The Digital Visual Interface (DVI) specification created by the [Digital Display Working Group](#) is designed to stimulate the widespread adoption of digital displays for high-performance desktop and mobile PCs. DVI provides the industry with a standardized display interface that, in turn, gives vendors a new way to display products, including simplified graphics subsystems and monitors, at lower cost.

Compatibility and Scalability

Removing analog graphics technology from the PC is the next logical step in the PC industry's legacy removal strategy, which is designed to make PCs simpler and easier to use. With DVI, the industry has a single standard interface that improves image quality by providing a digital signal from the PC's graphics subsystem to the display.

DVI uses an interface known as Transmission Minimized Differential Signaling (TMDS). It offers backward compatibility with existing standards while providing scalability in bandwidth and features. DVI supports two TMDS links, both with a transmission rate of 1.6 GB/sec.

With capabilities for copy protection, bidirectional communication, and selective refresh, DVI is projected to have a minimum life of 10 years.

DVI for LCDs...and CRTs

Although a DVI is often thought of as a flat-panel interface, significant cost savings can be realized when DVI is coupled with a digital-input CRT (see Table 1 below). Even with recent drops in the prices of flat panel displays, LCDs are likely to remain at least 50 percent to 65 percent more expensive than a CRT. The DVI standard presents the industry a significant opportunity to re-energize the CRT market by "going digital" and making the digital-input CRT a lower-cost alternative to the current generation of LCD panels.

Display Type	Cost	Performance	Viewing Angles	EMI	Weight/Size	Power Usage
CRT (today)	19" \approx \$350.00	UXGA, 1600x1200	180°	Not for magnetic areas	80 lbs.	\approx 100W
Digital Panel	18" \approx \$3500.00	SVGA, 1280x1024	120°	Suitable most areas	< 40lbs	< 6w

Table 1. Digital Panels and CRT Display Cost/Performance

Comparing Display Interfaces

As the Table 2 shows, DVI delivers digital signal from source to display to preserve signal integrity and the viewable image quality.

Display Type	Graphics Adapter	System output	Monitor Interface	Display Interface
Analog CRT	Digital-to-Analog	Analog	Analog	Analog
Analog Flat Panel	Digital-to-Analog	Analog	Analog-to-Digital	Digital
Digital CRT	Digital	Digital	Digital-to-Analog	Analog
Digital Flat Panel	Digital	Digital	Digital	Digital

Table 2. Signal Transmission from Source to Display

Improving Viewable Quality

The viewable quality of a digital CRT exceeds the quality of a traditional analog CRT by reducing the perceived fuzziness of characters and images. In today's PCs, the digital binary value of an image is converted by the graphics subsystem to the analog signal required by the CRT, degrading the signal quality. The analog signal is then transmitted at a slow rate over an inferior cable, which can further degrade the overall viewing quality.

By comparison, when a display adapter uses DVI, the image to be displayed does not go through this digital-analog conversion, preserving the original integrity of the digital signal until it reaches the display. What's more, the digital transfer rate of 1.6 Gb/sec. allows large quantities of display data to be moved to the display more rapidly.

Simpler and Better CRT Displays

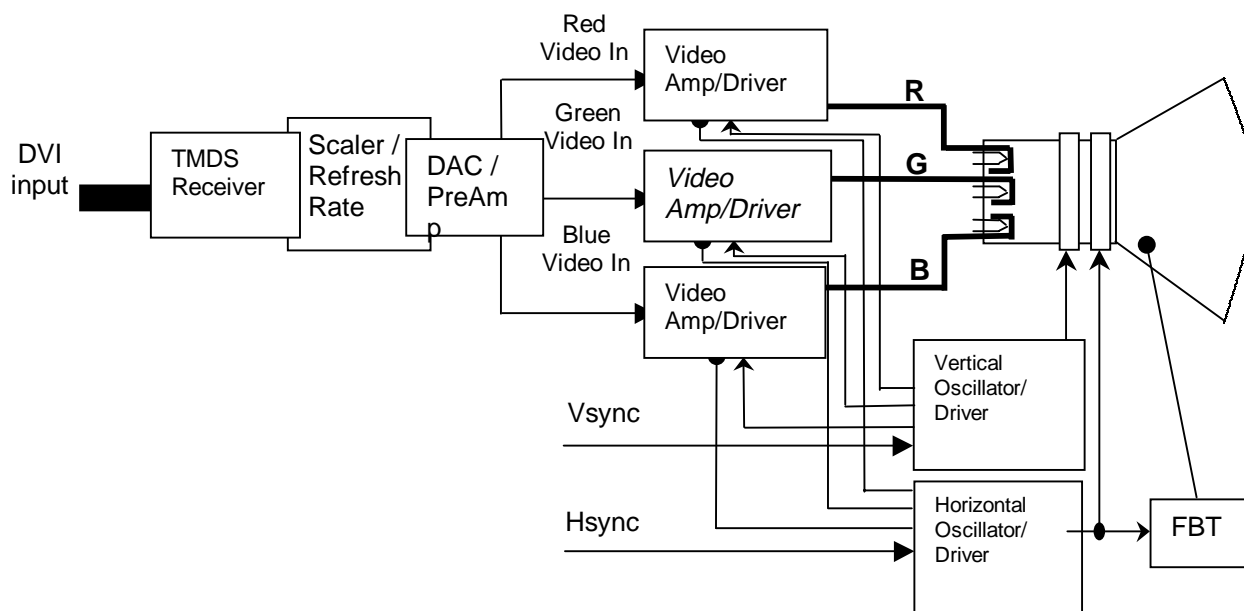
Current generation analog displays must support multiple refresh rates and resolutions, something that adds cost to both the display and the graphics subsystem. A digital CRT can have a fixed frequency and resolution like an LCD display and eliminate the need for multisync technology.

With DVI, screen refresh functionality can be part of the display itself. New data needs to be sent to the display only when changes to the data needs to be displayed. With this selective refresh interface, DVI can maintain the high refresh rates required to keep a CRT display ergonomically pleasing while avoiding an artificially high data rate between the graphics controller and the display.


A digital CRT can accommodate other beneficial features. Frame buffer memory and a digital/analog converter (DAC) are just two examples of added functionality that supports high performance while simplifying the PC-to-display combination.

The arrival of the DVI specification will help PC displays look better in terms of viewable image quality, features, and cost. For additional information about the DVI specification, visit the [Digital Display Working Group](#) Web site.

Analog CRT to Digital CRT: Block Diagram Examples



Example 1



Example 2

About the Author

Carol Jacobson is Intel's Initiative Technical Manager for the Digital Visual Interface initiative, where her responsibilities include managing a wide range of DVI-related technical, marketing, and industry-enabling activities. She has worked at Intel for over 16 years, during which time she has been involved in many different chassis and motherboard designs and enabling technology to the industry. Carol was the Intel Editor-in-Chief for PC99 with Microsoft and the industry. Prior to that she worked with the graphic controller community in delivering AGP.

Intel® Motherboards: Reliability and Support Provide System Integrators with a Competitive Edge

Chuck Fried
President and CEO
The Tech Shop

What happens when your product works so well it never breaks down? The Tech Shop, a systems integrator that uses only Intel® motherboards, has an answer: focus on technical services that give you a competitive advantage. Here's what to look for in a multi-pronged system of service, support, and training.

As a systems integrator (SI), when you're competing for the hardware accounts of the small business market, you're competing not only with hundreds of local dealers, but also with the world's top direct computer system companies. To advance, you can either play the game of the low-cost, high-volume dealer—and thereby expose already narrow margins to even tighter constraints—or you can rise above the crowd by differentiating yourself as a top-quality niche player.

The Tech Shop is a systems integrator based in Amherst, New York which has focused on reliability and strong vendor support in choosing its system motherboard supplier. Reliability translates into a minimal number of boards arriving “dead on arrival,” which ultimately means less reworks or customer returns. But product reliability is just part of the story; a supplier able to provide a wealth of support information and associated SI partner programs is also key to success.

For these reasons, The Tech Shop uses only Intel® motherboards when integrating systems for customers. Overall, the Intel brand gives companies like The Tech Shop the credibility we need to compete in the big-name clone market, and the quality assurance required to compete against low-cost dealers. The Intel Product Dealer Program provides SIs with the kind of useful information, training, and other support services that they can leverage to gain an edge in what is a highly competitive marketplace.

The Quality Challenge: Reliability Is Key

Before settling on the Intel® SE440BX-2 Motherboard, The Tech Shop sourced products from many different motherboard manufacturers. But the firm's rate of return was as high as 10 percent—meaning that for every \$1 million generated in system sales, \$100,000 typically would be tied up in return cycles for three to four weeks.

But these costs are a relatively small part of a much bigger picture. While hardware serves as the hook to catch a system customer, as much as 70 percent of the revenues that high-quality system integrators generate is typically earned from associated consulting services. If customers see a box in warranty repair, it severely hampers a system integrator's ability to market and ultimately sell those consulting services. In addition, every hour spent on warranty repair translates into lost opportunities. Instead of working on a repair, a technician or consultant could be building a new system, or generating fee-based revenues.

The Tech Shop chose Intel motherboards in part because Intel holds its component vendors and its own fabrication facilities to the highest reliability standards in the industry. In addition, Intel assimilates dealer feedback and suggestions into successive motherboard designs, yielding a stunningly low 2,000 defects per million—a rate five times better than the industry average, and equivalent to 0.2 percent customer-line-fallout rate. The record speaks for itself: The Tech Shop has not had a single motherboard-related product return or failure in the several years since it began integrating with Intel motherboards like the SE440BX-2.

Comprehensive Support for System Integrators

So what happens with a product that works so well that it never breaks down? In The Tech Shop's case, it means the kind of support that is required focuses less on fire drill-style repair and technical services than it does on other things that can help an SI gain a competitive advantage. This factor played a significant role in The Tech Shop's decision to standardize on Intel, which backs the quality and reliability of its motherboards with a multi-pronged system of service, support, and training.

Systems integrators working with Intel have the opportunity to attend product and technology training programs three times a year. In addition, they can receive extensive and up-to-date information on Intel's dedicated dealer Web site, where they can also access dealer news groups and leverage rapid and responsive communications via both e-mail and telephone. The Tech Shop frequently takes advantage of these resources, "cutting and pasting" Intel product literature, technical specifications, and information obtained from the dealer Web site for use in marketing to and communicating with clients.

Here's a closer look at what the Intel Product Dealer Program can provide:

- Technical instruction three times per year
- Product roadmaps under non-disclosure agreement
- Sales skills classes
- Q&A sessions
- Bi-monthly access to I.P.D. Net, including:
 - Technical information
 - Reference platforms
 - Advance notice on price moves
 - Weekly product availability reports
 - Exclusive promotions
 - Online training
 - Online merchandise ordering
 - Flyer Builder
 - Rebate payment status for your account
 - Ordering information

Compatibility: Up to the Test

Quality, reliability, support and training are just a few of the benefits that system integrators can reap by using Intel motherboards. In addition to taking advantage of these characteristics, neither The Tech Shop nor its customers has to spend valuable time dealing with interoperability or compatibility issues. That's because Intel tests all of its motherboards for BIOS functionality to ensure that the operating systems running on the motherboard are capable of using every component on the test system. The boards are then tested for compatibility with up to 350 commercially available and soon-to-be-released hardware, software, and networking products across the industry.

As The Tech Shop has discovered, building systems by standardizing on a motherboard solution from one supplier provides great efficiencies for the assembly line. The familiarity gained from using the Intel SE440BX-2 in all desktop configurations enables the company to integrate motherboards much more quickly. There's also a compelling and elegant simplicity involved in dealing with a full-line supplier. With but a single channel representative, tech support line, training program, and dedicated Web site for keeping abreast of technical issues and product news, The Tech Shop is able to focus on its core business, instead of a bulging Rolodex* comprised of all sorts of vendor protocols and support lines. And in the highly competitive world of system integration, that really matters.

About the Author

Chuck Fried founded The Tech Shop in 1992, and has served as the company's president and chief executive officer since its inception. Based in Amherst, New York, the Tech Shop offers a wide range of Internet, networking, and IT consulting services for small businesses in Western New York, and throughout the United States.

Prior to founding The Tech Shop, Fried ran a non-profit social services agency and held a variety of sales positions. He holds a Bachelor of Science degree in Business, with a concentration in finance, from the State University of New York (SUNY) at Buffalo.

For More Information

For more details on Intel Motherboards and The Tech Shop, please visit the following Web sites:

- [Intel® Motherboards](#) (channel information)
- [Intel® Motherboards](#) (technical information)
- [The Tech Shop](#)

Technology News Bytes

July 26

[Bluetooth\(tm\) Special Interest Group Announces Release of the Bluetooth 1.0 Specification](#)

The five founding companies of the Bluetooth™ Special Interest Group (SIG), Ericsson, IBM, Intel, Nokia and Toshiba, today announced the release of the Bluetooth 1.0 specification. With the completion of the Bluetooth 1.0 specification, developers from around the world can complete the design of products equipped with Bluetooth technology and prepare for product qualification and interoperability testing.

July 20

[NGIO Forum Announces Release of 1.0 Specification](#)

The NGIO (Next Generation I/O) Forum® today announced the immediate release of the NGIO 1.0 specification. The channel-based, switched-fabric architecture provides a new high speed, low cost, reliable and scalable data transfer method. It is designed to allow future servers to uncompromisingly meet customer requirements for high speed transfer of data between servers, storage, networks and other attach points.

July 20

[Intel Introduces File Sharing Appliance for Small Business Customers](#)

Intel introduced the Intel® InBusiness™ Storage Station, a network appliance designed to help small businesses increase productivity through collaborative file sharing, protecting of critical files, and significantly increasing electronic file storage capacity.

July 15

[Intel Enters Broadband Market to Enable Faster Internet Access for PC Users](#)

Intel announced plans to deliver broadband connectivity devices aimed at providing residential and small business users faster access to the Internet. An ADSL technology licensing agreement with Cisco Systems, Inc. is Intel's first step to deliver broadband access technology. Intel's entrance into the emerging broadband marketplace marks its intent to make high-speed access--typically 25 times faster than the fastest analog modems--more affordable and widely available to end users.

July 14

[Intel Forms a New Business Unit and Aligns Existing Groups around Internet Opportunities](#)

Intel announced the formation of a new business unit along with related changes within other business groups designed to increase Intel's focus on opportunities surrounding the Internet. Intel has formed the new Communications Products Group to drive the company's systems initiatives in networking and communications, combining related hardware, software and support capabilities into one organization. The new group will be headed by Intel vice president John H. F. Miner, formerly general manager of the Enterprise Server Group (ESG).

July 13

[Intel Completes Merger with Dialogic](#)

Intel announced the completion of its previously announced merger with Dialogic Corporation. The merger expands Intel's standard-high-volume (SHV) server business in the multibillion-dollar networking and telecommunications market segment by providing industry vendors with standards-based hardware and software building blocks for integrated voice and data networks.

July 7

[Intel to Acquire Softcom](#)

Intel announced it has entered into a definitive agreement to acquire privately held Softcom Microsystems, Inc. in an all cash transaction. Softcom develops and markets semiconductor products for original equipment manufacturers (OEMs) in the networking and communications market segments. The company's high performance components are designed for networking gear (access devices, routers, and switches) used to direct voice and data across the Internet as well as traditional enterprise networks.

June 30**[Leading Corporations Move to Intel and Windows NT-Based Workstations](#)**

Some of the world's leading companies from a wide range of industry sectors are announcing their move to Intel® Corporation and Microsoft® Windows NT® operating system-based solutions at the Workstation Leadership Forum (WLF) '99. Companies such as BankBoston, Ellerbe Becket, Enron Corp., Lear Corp., Navistar International Corp., Sony Pictures Imageworks and Xerox Corp. are among the corporations highlighted at the second annual workstation event, hosted by Intel and Microsoft Corp.

June 30**[Intel and SGI Announce Agreement to Optimize OpenGL API for Current and Future Intel Architecture Platforms](#)**

At Intel's Workstation Leadership Forum, Intel and SGI today announced plans to work together to optimize the OpenGL® API for advanced 3-D graphics and data visualization workstations based on Intel Architecture (IA). This joint effort is expected to dramatically increase the performance and visualization capabilities of IA-based workstations, utilizing OpenGL API graphics, from entry-level to the highest-end platforms.

June 28**[Intel Launches Worldwide Channel Initiative for Internet Service Providers](#)**

Intel introduced the Intel Internet Service Provider Program, a worldwide channel initiative that will provide Internet Service Providers (ISPs), Application Service Providers (ASPs), and value-added resellers (VARs) with a variety of Internet-tailored products and technologies sold by a dedicated channel sales and support team. Aimed at the estimated 14,000 ISPs and ASPs worldwide, the Intel Internet Service Provider Program is a response to the rapidly changing nature of the Internet economy.

Industry Events

Industry Events for August

Herring on Hollywood

August 1-3
Century Plaza Hotel
Los Angeles, CA

Digital technology has made it easier for the global entertainment industry to create entertainment and to more effectively push its content to a worldwide audience. But at the same time, digital technologies—especially the Internet—have radically altered distribution models; created tensions between the entertainment, computer, and telecom industries; facilitated new competition in brand recognition; made copyright laws almost irrelevant; and thrown retail margins up for grabs. Business models are changing, and the entertainment industry has been turned upside down. Find out what it all means at Herring on Hollywood's two-day conference, with the theme "Convergence: Deal with It." Intel is a Headline Host.

Intel's Wired for Management Summit '99

August 2, 1999
Santa Clara Convention Center
Santa Clara, CA

If you're an IT professional, manager or executive with responsibility for managing PCs, Intel's Wired for Management Summit '99 can show you how Wired for Management (WfM) enhanced products can improve management of complex computing environments and reduce Total Cost of Ownership (TCO). You'll see case histories from large corporations using WfM-enhanced solutions and hear how these companies are using real-world products to increase their strategic focus and improve their organizations' competitive agility. A Solutions Showcase will offer hands-on demonstrations by leading system management companies. The Summit will feature an exciting, informative live event including a keynote by an industry analyst with the Gartner Group.

SIGGRAPH 99

August 8-13, 1999
Los Angeles Convention Center
Los Angeles, CA

ACM's annual gathering for its Special Interest Group on Graphics has become the world's epicenter of computer graphics and interactive technologies, with hundreds of companies offering thousand of products and services for the professional digital content creator.

With the theme "Shorten the Distance Between Thinking It and Seeing It," Intel is taking the workstation to a new dimension at Siggraph. A stereoscopic 3D theater presentation will be this year's highlight in the Intel booth (#701). All content is original and has been created on Intel® Architecture-based workstations. Segments from leading studios (Sony Imageworks, RezN8, Digital Depth, Blur, Threshold, and Digital Domain) will be featured. The booth will also offer hands-on workshops, a digital content creation studio, a workstation technology showcase, and a display wall focusing on compatible graphics add-ons from IHVs.

The CIO 100 Symposium & Awards

August 15-18
Hotel del Coronado
San Diego, CA

The CIO 100 Symposium & Awards is a three-day program focusing on Leadership and Innovation for the Future of the Enterprise. Through interactive presentations and discussions, the Symposium brings to life and honors the achievements of 100 industry-leading enterprises. Intel is a partner in this event and its Enterprise Server Group will be participating.

Hot Chips Conference

August 15-17

Stanford University

Palo Alto, CA

Since its beginning in 1989 Hot Chips has become one of the leading conferences in the microprocessor/microcomputer field. Its emphasis is on real products and real technology. Hot Chips is sponsored by the IEEE Computer Society, Technical Committee on Microprocessors and Microcomputers.

Seybold San Francisco 21st Century Publishing

August 30- Sept 3

San Francisco, CA

Each day of SSF has something to help people in electronic publishing do their jobs better, make more money, or understand how the business environment is going to change. The focus will be on the open exploration of the key issues and the introduction of new concepts and technologies for the electronic publishing process, with clues to the direction the industry will be taking. Tracks include: the state of Internet design and technology, optimal server architecture, connectivity, mixed-platform environments, and digital imaging.

Intel Developer Forum (IDF)

August 31-Sept 2

Palm Springs Convention Center

Palm Springs, CA

Three full of in-depth presentations, demonstrations and dialogue with Intel's chief technology architects, plus third party luminaries will give you a head start for designing platform for the new millennium. The Intel Developer Forum covers today's implementation details and tomorrow's technology roadmaps to keep you up-to-date on the latest desktop, workstation, server, and mobile platforms and embedded technologies including:

- IA-64—emerging Internet server technology
- Enabling trusted Internet access and e-commerce—year 2000 security roadmap
- Intranet/enterprise clients and servers
- Software tools for authoring rich Internet content and applications
- Easy-to-use PCs and home networking—bringing Internet access to every room in the home
- Next-generation StrongARM[®] solutions—Internet applications and digital companions
- Improving and extending Internet connectivity—broadband, wireless
- USB 2.0—higher bandwidth connectivity for PC peripherals

Keynotes will be delivered by Intel's top executives and will provide perspectives on how these technologies are shaping the PC Industry.

Industry Events for September

System Builder Summit

September 9

Boca Raton Resort and Club

Boca Raton, FL

This is an exclusive forum that brings together the powerhouses in the white-box market with leading technology vendors. In 1999, white-boxes will account for 42% of the sales of new PCs. The System Builder Summit hosts the top 1% of the leading system builders who represent significant purchasing power. System builders, who attend as guests of the event, discuss critical channel issues with each other and with senior executives from manufacturing, software, distribution and service organizations.

Boon Chye Ooi, Vice President, Intel Desktop Products Group, and General Manager, OEM Platform Solutions Division, keynotes at SBS.

NetWorld+Interop 99

September 13-17

Georgia World Congress Center

Atlanta, GA

Mark Christensen, Vice President and General Manager, Intel Network Communications Group, keynotes Tuesday, September 14 from 5:00 to 6:00 p.m.

Embedded Systems Conference

September 27-30

San Jose Convention Center

San Jose, CA

Intel's Computing Enhancement Group and its Applied Computing products and technologies will be represented in a booth. Vice President and General Manager Ron Smith will be a featured speaker.

Note: ARM and StrongARM are trademarks of Advanced RISC Machines, Ltd. Third-party brands and names are the property of their respective owners.

Industry Events for October

Embedded Internet Workshop

October 1

San Jose, CA

Tom Franz, Vice President and General Manager, Intel Embedded Microcomputer Division, is the keynote speaker who opens the conference.

Fall Internet World '99

October 4-8

New York, NY

Whether you're a webmaster evaluating E-commerce software, an ISP rethinking business strategy, a marketing executive looking for proven online opportunities, or an IT professional who needs to secure an enterprise network, Internet World has something for you. Billed as "The World's Largest Event for E-Business and Internet Technology," Fall Internet World 99 rivals the best stand-alone, vertical conferences in the industry for both content and emerging leadership. From advanced technology workshops to introductory sessions, attendees can choose from a range of programs covering every aspect of conducting business on the Internet.

Intel is the corporate sponsor of the conference. Intel President and CEO Craig Barrett gives the opening keynote 9:00-9:45 a.m. Wednesday, October 6.

Gartner Group U.S. Symposium/ITxpo 99

October 12

Lake Buena Vista, FL

Symposium/ITxpo 99 will address the entire spectrum of IT with the multiple needs of business in mind—now and in the future. If you plan, manage, strategize, guide, purchase, market, direct or use IT, Symposium/ITxpo is a crucial strategic tool.

The Symposium will conduct a Mastermind Interview with Intel President and CEO Craig Barrett. In this interview Gartner will ask Dr. Barrett about Intel's view of global connected computing, advanced networking, distributed vs. central computing, and the future of the microprocessor. Gartner analysts will question him on Intel's strength, strategies, partnerships, and competition.

United States Telephone Association 102nd Annual Convention & Exhibition

October 17-19

Hyatt Regency Hotel

San Francisco, CA

The USTA Annual Convention brings together senior executives from all segments of the telecommunications industry. Telecom executives attend the annual event to network with other industry professionals, participate in high-level discussions about critical industry issues and to gain valuable insights that will help them grow their businesses. Over 1,000 participants are expected to attend the Convention, representing every area of the communications industry. This year's theme is "Communications Convergence: Focus on Customer Solutions."

The opening keynote speakers will be Intel co-founder and chairman Andy Grove and FCC Chairman William Kennard.

Internet Commerce Expo (ICe)

October 18-21

Moscone Convention Center

San Francisco, CA

Internet Commerce Expo is the leading eBusiness trade show and conference focusing on Internet Commerce, intranets, extranets and Web-based business solutions. ICe is specifically targeted at business users and IT professionals that are buying Web-based technologies and applications to advance productivity and profitability gains in the eBusiness enterprise.

Intel will be exhibiting at ICe. Paul Otellini, executive vice president and general manager, Intel Architecture Business Group, will present a keynote address on October 20, 9:30 – 10:30 a.m.

DV Expo 99

October 19-22

Long Beach Convention Center

Long Beach, CA

DV Expo is the place to meet peers and experts, learn time-saving techniques, and get the low-down on the latest tools and technologies. The Digital Video Conference & Expo offers four intense days of in-depth training on digital video technology and business. It puts attendees on the fast track for learning the latest tips and tricks and evaluating the best product solutions for digital video projects. Topics include production, editing, motion graphics, 3D, visual effects, corporate video, compression, DVD and Web, and more. Intel is an exhibitor.

ISPCON Fall '99

October 26-28

San Jose Convention Center

San Jose, CA

ISPCON Fall '99 will gather thousands of Internet Service Providers, CLECs, ICSPs, ITSPs, IXCs, Cable Operators, Next Gen Telcos, Telecom Resellers, Satellite/Wireless Companies, and Web Hosting companies for an intensely educational event focused on the business of providing Internet access.

Intel is an exhibitor and corporate sponsor. Communications Products Group vice president and general manager John Miner will be speaking during the conference.

1999 NGIO Forum's Developers' Conference

September 28-30

Newport Beach, CA

Intel is a founding member and a steering committee member of the NGIO Forum.

—End of Platform Solutions Issue 22—